AD-A145 612 MEMORY-BASED EXPERT SYSTEMS(U) YALE UNIV NEW HAVEN CT DEPT OF COMPUTER SCIENCE R C SCHANK AUG 84 AFOSR-TR-84-0814 F49620-82-K-0010 1/1

UNCLASSIFIED F/G 9/2 NL

END

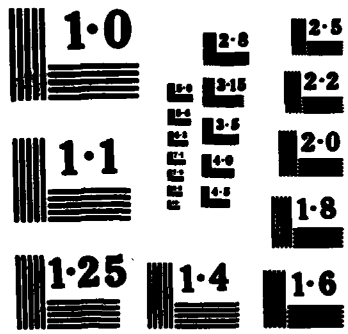
1·0
1·1
1·25
1·4
1·6
1·8
2·0
2·2
2·5

6a. NAME OF PERFORMING ORGANIZATION Yale University	6b. OFFICE SYMBOL *(If applicable)*	7a. NAME OF MONITORING ORGANIZATION Air Force Office of Scientific Research
6c. ADDRESS *(City, State and ZIP Code)* Computer Science Department, P.O. Box 2158, Yale Station, New Haven CT 06520		7b. ADDRESS *(City, State and ZIP Code)* Directorate of Mathematical & Information Sciences, Bolling AFB DC 20332
8a. NAME OF FUNDING/SPONSORING ORGANIZATION AFOSR	8b. OFFICE SYMBOL *(If applicable)* NM	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER F49620-82-K-0010
8c. ADDRESS *(City, State and ZIP Code)* Bolling AFB DC 20332		10. SOURCE OF FUNDING NOS.

PROGRAM ELEMENT NO.	PROJECT NO.	TASK NO.	WORK UNIT NO.
61102F	2304	A2	

11. TITLE *(Include Security Classification)*
MEMORY-BASED EXPERT SYSTEMS

12. PERSONAL AUTHOR(S)
Roger C. Schank

13a. TYPE OF REPORT	13b. TIME COVERED	14. DATE OF REPORT *(Yr., Mo., Day)*	15. PAGE COUNT
Interim	FROM 15/1/83 TO 14/1/84	AUG 84	8

16. SUPPLEMENTARY NOTATION

17. COSATI CODES			18. SUBJECT TERMS *(Continue on reverse if necessary and identify by block number)*
FIELD	GROUP	SUB. GR.	

DTIC ELECTE SEP 1 9 1984 S E

ITEM #19, ABSTRACT, CONTINUED: ~~compue~~ come up with an answer. This second level of reasoning should more closely resemble human problem solving behavior when people are confronted with novel situations. (3) A cornerstone to this method is automatic learning. The system's memory of experiences will be changed and augmented by each additional case that is presented. The system will remember the problems that it has encountered and use that information to solve future problems. These three principles of the memory-based expert systems model are being tested in several related projects.

Accession For	
NTIS GRA&I	☒
DTIC TAB	☐
Unannounced	☐
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

DTIC
COPY
INSPECTED
3

AFOSR-TR- 84-0814

MEMORY-BASED EXPERT SYSTEMS
INTERIM TECHNICAL REPORT
AFOSR CONTRACT F49620-82-K-0010
15 JANUARY 1983 - 14 JANUARY 1984

Roger C. Schank
Department of Computer Science
Yale University
New Haven CT 06520

Approved for public release;
distribution unlimited.

Roger C. Schank
Yale Department of Computer Science

Memory-Based Expert Systems
Contract No. F49620-82-K-0010

Chief, Technical Information Division

In the past dozen years, one of the most visible developments in artificial intelligence research has been the emergence of rule-based expert systems. These programs have been applied to more and more domains that require extensive knowledge for very specific and critical tasks, including medical diagnosis, hardware troubleshooting, and geological exploration. These systems have come under severe scrutiny and, despite their capable performance in certain task domains, seem to be subject to several critical problems:

1. It is extremely difficult and time-consuming to construct the knowledge base.
2. The programs are unable to deal with problems that are not explicitly covered by the rule-base.
3. Most additions to the programs require programmer intervention.

We are developing a model of expertise that more closely resembles the way in which humans become experts, namely, through *experience.* We assume that the rule-base is not the primary repository of knowledge, but rather rules are derived from experience. Our model addresses the three problems given above as follows.

1. The knowledge-base is derived primarily from the enumeration of specific cases or experiences. We have found that a human expert is much more capable of recalling experiences than articulating internal rules. We suggest that the reason for this difference is that the human expert may not in fact be using rules in the first place.
2. As problems are presented to the system for which no specific case or rule can match exactly, the system can reason from more general similarities to come up with an answer. This second level of reasoning should more closely resembles human problem solving behavior when people are confronted with novel situations.
3. A cornerstone to this method is automatic learning. The system's memory of experiences will be changed and augmented by each additional case that is presented. The system will remember the problems that it has encountered and use that information to solve future problems.

These three principles of our memory-based expert systems model are being tested in several related projects.

The ALFRED Economics Learning Project

The ALFRED economics learning project has been through three phases since the last report:

- knowledge representation and reasoning style
- text argument interpreter and primitive chain builder
- large-scale intensive data collection

Phase 1: knowledge representation

This research is reported in *Knowledge reorganization and reasoning style,* given in the Publications below. In the previous year, the economics project had developed a particular simple economic reasoning system, whose primary purpose was to exercise our representation of economic knowledge and test the rules of economic inference. The system was not, however, a particularly good model of novice economic reasoning, and did not reflect any theory of how economic expertise developes.

Using several dozen protocols gathered by William Salter, a psychology graduate student, for his Ph.D. thesis, we developed two distinct models of economic reasoning. One, the expert model, is based on the model from the previous year. It reasons about economic *quantities* and their connections, e.g., "As interest rates increase, capital investment is stifled." This is the kind of reasoning most commonly found in the *Wall Street Journal* and other sources of economic analyses.

The second model, the novice. is based on Carbonell's POLITICS model. It reasons about economic *actors*, their goals, plans, and behavior, e.g., "As taxes decrease, workers won't need to strike for higher salaries."

The important observation about the novice model is that it is capable of *understanding* the arguments produced by expert reasoning, but is too weak to generate most of them. The novice model knows how to search its knowledge base to find out how economic events, such as tax cuts, might affect the novice and other actors, but it does not know how to search the knowledge base to calculate the net effect of some event on some economic quantity, such as inflation or the Federal deficit.

Since the novice can understand expert arguments, however, the novice model can learn to argue better, but changing its search rules so that they gradually begin to follow the same paths that the experts do.

Phase 2: text interpreter

The only way to test and develop a theory of economics expertise acquisition is to have real examples to learn from. Although we have experimented with simplified English inputs and with hand-analyzed inputs, any serious model has to deal with the real input texts that people get, for several reasons:

- If the learning system is to stand on its own, it must be able to handle real texts.
- People learn from texts they only partially understand. Any automatic learning system must have same ability.
- Language understanding are both memory-based and must be able to work together in harmony.

The ALFRED argument interpreter, written by James Spohrer, uses a parser based on on Lebowitz' Integrated Partial Parser. Parsing is treated more as the identification of stored memory structures than as the construction of meaning representations. It currently handles only a few examples of real texts (but see "Phase 3"), and only uses expert-like quantity-based reasoning.

The following is an example of one of the texts that it analyses, an argument given by Lester Thurow in a *Newsweek* article, followed by its interpretation, converted into readable form by a very simple data structure describer. Note that the analysis keeps track not only of Thurow's argument, but also of what Thurow says is the President's argument. Also notice that the analyzer handles not only the causal chaining, but also the use of an example as an argument support.

```
ARGUMENT TEXT
  The President says that if Americans are given a 25 per cent
  across the board tax cut, they will suddenly become large
  savers. Last year the average american saved 5.6% of his income.
  Americans have never been large savers.  The most likely
  hypothesis is that they will consume most of their tax cut.
  But without extra personal savings to balance the extra investment
  stimulated with the large business tax cuts, interest rates
  can only rise.

ARGUMENT REPRESENTATION

  AUTHOR: THUROW
  KIND OF ARGUMENT: EXAMPLE CAUSAL

  PREMISE: increasing IND:INCOME and decreasing GOV:TAX-RATE

  CONCLUSION: increasing BNK:INTEREST-RATE

SUPPORT CHAIN
```

```
increasing IND:INCOME causes decreasing IND:SAVINGS
  because...
    there is an example where increasing IND:INCOME
    causes increasing IND:CASH
    causes decreasing IND:SAVINGS

decreasing IND:SAVINGS causes increasing IND:SPENDING
  because... decreasing IND:SAVINGS
    causes increasing IND:CASH
    causes increasing IND:SPENDING

decreasing GOV:TAX-RATE causes increasing IND:SPENDING
  because... decreasing GOV:TAX-RATE
    causes decreasing IND:TAX-RATE
    causes decreasing IND:PAY-TAXES
    causes increasing IND:CASH
    causes increasing IND:SPENDING

increasing IND:SPENDING causes decreasing IND:SAVINGS
  because... increasing IND:SPENDING
    causes decreasing IND:CASH
    causes decreasing IND:SAVINGS

decreasing GOV:TAX-RATE causes increasing BUS:INVEST
  because... decreasing GOV:TAX-RATE
    causes decreasing BUS:TAX-RATE
    causes increasing BUS:E-ROI
    causes increasing BUS:D-INVEST
    causes increasing BUS:BORROW
    causes increasing BUS:INVEST

increasing BUS:INVEST causes increasing BNK:INTEREST-RATE
  because... increasing BUS:INVEST
    causes increasing BUS:BORROW
    causes increasing ECONO:BORROW
    causes increasing ECONO:CREDIT-DEMAND
    causes increasing BNK:INTEREST-RATE

decreasing IND:SAVINGS causes increasing BNK:INTEREST-RATE
  because..., decreasing IND:SAVINGS
    causes decreasing BNK:CASH
    causes decreasing BNK:D-LOAN
    causes increasing BNK:INTEREST-RATE
 }
```

MENTIONED SUBARGUMENT

```
ARGUMENT REPRESENTATION

  AUTHOR: PRESIDENT
  KIND OF ARGUMENT: CAUSAL

  PREMISE: decreasing GOV:TAX-RATE

  CONCLUSION: increasing IND:SAVINGS

SUPPORT CHAIN

decreasing GOV:TAX-RATE causes increasing IND:SAVINGS
  because... decreasing GOV:TAX-RATE
    causes decreasing IND:TAX-RATE
    causes decreasing IND:PAY-TAXES
    causes increasing IND:CASH
    causes increasing IND:SAVINGS
```

Phase 3: data collection

We have no intention of "solving" the natural language problem. We do need to solve the problem of natural language economic arguments however. This means that we need to gather data on what kinds of linguistic constructions, what kinds of arguments, what kinds of facts, and so on, occur commonly in economic arguments. Charles Martin has found and converted a number of relevant texts into machine-readable form, and constructed data analysis routines appropriate to our particular needs. This includes identifying domain-related words ('inflation', 'recession', etc.), common phrases, and basic conceptual categories. The preliminary results are:

Number of economics texts: 29

Number of lines: 2732

Average number of lines / text: ~94

Total number of words: 16678

Total vocabulary: 2591 words

1388 unique words (single occurrences)
~1180 domain-specific words identified
~1440 preliminary phrases identified
~70 preliminary conceptual categories identified

Preliminary analysis also shows that the ability to add new vocabulary items very easily will be crucial, since even after 10000 words of text have been processed, an average of 15% of the input still involves new words. Some of these are proper names and inflections of existing words, of course, but economic arguments are written for a literate audience with a large vocabulary.

Learning in a complex real-world domain

Learning research today must focus on real world domains, in which the complexity of the domains themselves forces researchers to confront the most difficult problems a theory must face. Early research on learning concentrated on tasks such as solving simple puzzles, such as the Towers of Hanoi, or learning very simple concepts, such as what is an arch. This project attempts to construct a model of a complex domain by making observations of the domain over time. The model which is constructed is used to predict future events, and when these predictions fail, hypotheses are constructed to explain the failure. These hypotheses then become part of the updated model of the domain. Because the domain is complex, the construction of new hypotheses is a potentially explosive process: the number of explanations for any event is extremely large. However, it has been found that causal knowledge of the domain greatly constrains the hypothesis generation process. When a failure occurs, knowledge of what events and features of the domain can cause other events allows a program to pinpoint more precisely the source of the failure, and therefore to construct fewer, more relevant hypotheses than it would otherwise.

A program has been built which implements these ideas in a complex, real-world domain. Performance of the program demonstrably improves over time and outperforms many human experts. Current work is examining the possible use of more detailed domain knowledge to focus the explanation process even further, including the construction of a physical model of the domain, and the addition of higher level structures for organizing the memory of past episodes.

Legal Reasoning: JUDGE

The practice of law has institutionalized the notion of reasoning from prior cases. It is therefore a most appropriate domain for our experience-based approach to expert systems.

JUDGE is a case-based reasoning project which intends to model a court judge's ability to decide on sentences for criminal cases. By *case-based* we mean that the system will use its experiences from similar previous cases to decide on an appropriate sentence for a newly input criminal situation. We are working with judges in New Haven to gather data for the program.

Some of the issues associated with this task include: assessing the likelihood that the criminal will repeat his crime, given the facts of the current case; his plea and his prior record;

accounting for social responsibility in sentencing, such as considering what society demands, and how the victims have reacted to the case; determining a sentencing plan, including the type of punishment, if any, and the duration, based on state law and on the other mentioned considerations in the case.

Further work on this system will involve formulating indexing, retrieval and application strategies for rules of justification so that chains of rules used in previous cases can be extracted and then modified for new cases, rather than having the system search through an entire rule base for guidance. That is, once the system has made a decision about the wrongness of a set of actions, any new case which involves a similar kind of wrongness should make use of the decisions from the old case. The system should either use the reasoning chain as is or it should modify it so that it better accounts for the facts of the new case. Our intent is to avoid looking for applicable rules to form a reasoning chain from scratch for each input case.

Explanation and reasoning

We have recently started work on developing a theory of *explanation* that can be applied to computer systems. That is, an intelligent program should be able to explain both its own actions and the actions of others. Very often this explanation process will involve reference to specific experiences and thus is a natural application of the memory-based theory of expertise which we are exploring.

When we first began to study reminding, it seemed clear that explanations held the key to reminding. Whenever two phenomena were related in a reminding experience, we found an explanation in common. Both events, the "remindee" and the "remindand" were analogous in the same way, usually an expectation failure. The explanation of the analogues was the link between them. Thus, in an important way, explanation and learning were already linked. To make sense of, and correct an expectation failure, explanations were made! These explanations were indices to the events they explained and, as such, could cause remindings. The result of all this was a "corrected expectation" or in some cases, a reorganized set of expectations. Thus the link went ---> failure-expectation - reminding - generalization - learning (modifiation of cause of failure).

Now, as we have progressed in our work on reminding and memory organization, it's becoming clearer that we missed the mark somewhat. The above chain is right enough, learning does occur in this way. But, the focus is wrong. Previously we focussed on the value of reminding and the correction of expectation failure, whereas we should have been focussing on explanations themselves. Here's why:

Intelligent human beings seek to understand the world around them. They would like to understand the people with whom they interact to the extent of knowing what they may do and why they may do it. They seek to understand the intuitions that they deal with. They want to know how to treat the rules that these institutions set up and how the institution will treat them. They also want to know how the physical world around them behaves. They want to know why machines behave the way they do and how physical objects and forces can best be dealt with.

To summarize: people want to understand the world, personally, socially, and physically. This understanding and analysis of the world involves explanation. We believe that our advanced artificial intelligence programs must have the ability to explain themselves and the world around them.

Publications and talks

Bain, B. "Assignment of Responsibility in Ethical Judgements." Paper presented at First Annual Workshop on Theoretical Issues in Conceptual Information Processing, Atlanta, GA., March 1984.

Riesbeck, C. "Knowledge Reorganization and Reasoning Style." Research Report #270, Computer Science Department, Yale University, June, 1983. To appear in *International Journal of Man-Machine Studies*, **19**, 1983.

Riesbeck, C. "Parsing -- Is it interesting anymore?" Member of panel at First Annual Workshop on Theoretical Issues in Conceptual Information Processing, Atlanta, GA., March 1984.

Riesbeck, C. Participant at the Systems Control/Artificial Intelligence meeting, Air Force Academy, Colorado Springs, Colorado, June 28-29, 1983.

Salberg, S. "Generating Hypotheses to Explain Prediction Failures", Paper presented at AAAI-83, Washington, DC. August, 1983.

Schank, R. "Learning, Explanation, and a Little History." Invited talk at First Annual Workshop on Theoretical Issues in Conceptual Information Processing, Atlanta, GA., March 1984.

END

FILMED

10-84

DTIC